



Building the Next Generation of Parallel Applications

Michael A. Heroux
Scalable Algorithms Department
Sandia National Laboratories, USA

A Brief Personal Computing History



1988 - 1997

```
CMIC$ DO ALL VECTOR IF (N .GT. 800)
CMIC$1  SHARED(BETA, N, Y, Z)
CMIC$2  PRIVATE(I)
CDIR$ IVDEP
    do 15 i = 1, n
        z(i) = beta * y(i)
15  continue
endif
```



1993 - 2008

```
#include <mpi.h>
int main(int argc, char *argv[]) {
// Initialize MPI
    MPI_Init(&argc,&argv);
    int rank, size;
    MPI_Comm_rank(MPI_COMM_WORLD, &rank);
    MPI_Comm_size(MPI_COMM_WORLD, &size);
```



2008 - Present



Unification and composition:
-Vectorization
-Threading
- Multiprocessing

```
#include <mpi.h>
#include <omp.h>
int main(int argc, char *argv[]) {
// Initialize MPI
MPI_Init(&argc,&argv);
int rank, size;
MPI_Comm_rank(MPI_COMM_WORLD, &rank);
MPI_Comm_size(MPI_COMM_WORLD, &size);
...
#pragma omp parallel
{
    double localasum = 0.0;
#pragma omp for
    for (int j=0; j< MyLength_; j++) localasum += std::abs(from[j]);
#pragma omp critical
    asum += localasum;
}
```

```
#include <thrust/host_vector.h>
#include <thrust/device_vector.h>

thrust::device_vector<int> vd(10, 1);
thrust::host_vector<int> vh(10,1);
```



Quiz (True or False)

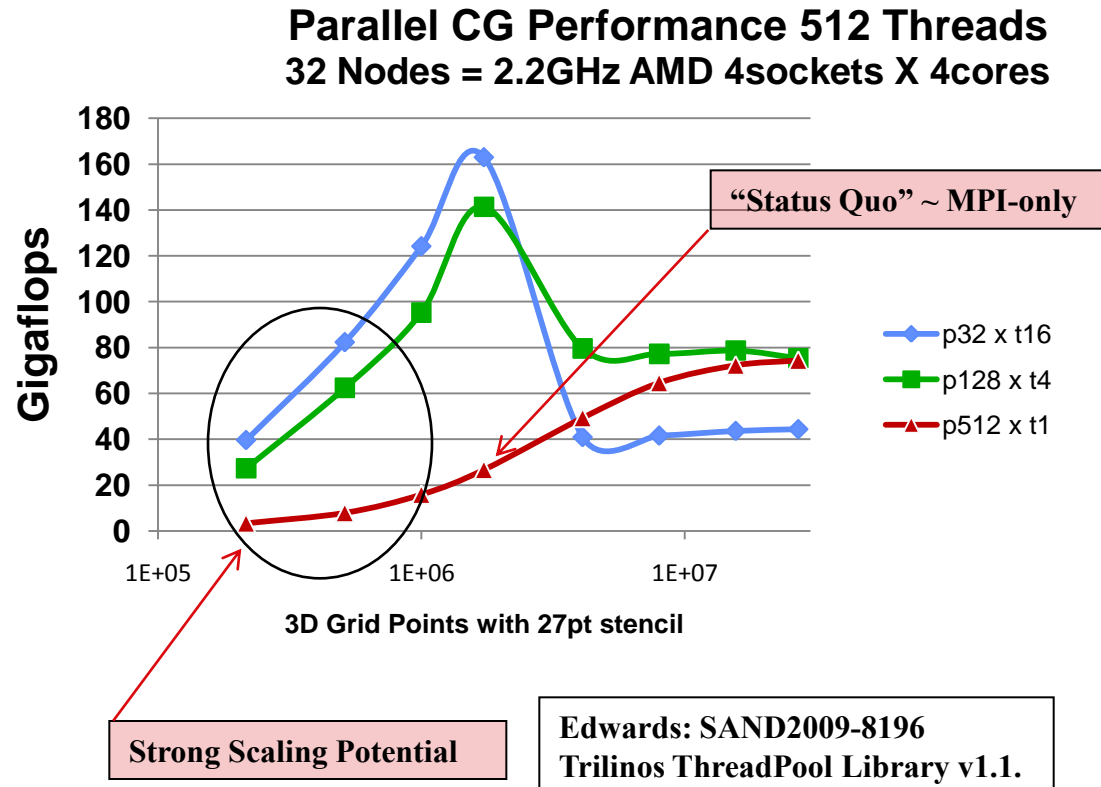
1. MPI-only has the best parallel performance.
2. Future parallel applications will not have MPI_Init().
3. All future programmers will need to write parallel code.
4. Use of “markup”, e.g., OpenMP pragmas, is the least intrusive approach to parallelizing a code.
5. DRY is not possible across CPUs and GPUs
6. GPUs are a harbinger of CPU things to come.
7. Checkpoint/Restart will be sufficient for scalable resilience.
8. Resilience will be built into algorithms.
9. MPI-only and MPI+X can coexist in the same application.
10. Kernels will be different in the future.

Basic Exascale Concerns: Trends, Manycore

- Stein's Law: *If a trend cannot continue, it will stop.*

Herbert Stein, chairman of the Council of Economic Advisers under Nixon and Ford.

- Trends at risk:
 - Power.
 - Single core performance.
 - Node count.
 - Memory size & BW.
 - Concurrency expression in existing Programming Models.



One outcome: Greatly increased interest in OpenMP



Implications

- MPI-Only is not sufficient, except ... much of the time.
- Near-to-medium term:
 - MPI+[OMP|TBB|Pthreads|CUDA|OCL|MPI]
 - Long term, too?
- Long- term:
 - Something hierarchical, global in scope.
- Conjecture:
 - Data-intensive apps need non-SPDM model.
 - Will develop new programming model/env.
 - Rest of apps will adopt over time.
 - Time span: 20 years.



What Can we Do Right Now?

- Study why MPI was successful.
- Study new parallel landscape.
- Try to cultivate an approach similar to MPI.



MPI Impresssions

MPI: It Hurts So Good

Dan Reed, Microsoft

Workshop on the Road Map for the
Revitalization of High End
Computing
June 16-18, 2003

• Observations

- “assembly language” of parallel computing
- lowest common denominator
 - portable across architectures
 - system independent
 - easy to learn
- upfront effort required

• C++

So What Would Life Be Like Without MPI?

$$F(n) = \begin{cases} 0 & n=0 \\ 1 & n=1 \\ F(n-1) + F(n-2) & n>1 \end{cases}$$

```
Serial C
long fib_serial(long n)
{
    if (n < 2) return n;
    return fib_serial(n-1) + fib_serial(n-2);
}
```

```
OpenMP 3.0
long fib_parallel(long n)
{
    long x, y;
    if (n < 2) return n;
    #pragma omp task default(none) shared(x,n)
    {
        x = fib_parallel(n-2);
    }
    y = fib_parallel(n-1);
    #pragma omp taskwait
    return (x+y);
}
```

```
Cilk++
long fib_parallel(long n)
{
    long x, y;
    if (n < 2) return n;
    x = cilk_spawn fib_parallel(n-2);
    y = fib_parallel(n-1);
    cilk_sync;
    return (x+y);
}
```

Tim Stitts, CSCS
SOS14 Talk
March 2010

```
def fib_ser
{
    var x,y: n;
    if (n < 2) then
        cobegin {
            x=fib_serial(n-1);
            y=fib_serial(n-2);
        }
    return x+y;
}
```

```
fib_parallel(n: Z): Z requires
if n < 2 then n else fib_pa
```



“MPI is often considered the
“portable assembly language” of
parallel computing, ...”

Brad Chamberlain, Cray, 2000.

Looking Forward to a New Age of Large-Scale Parallel
Programming and the Demise of MPI
...hopes and dreams of an HPC educator

[illegible]

```

do i=0,n2-1
do il=i,n1
buff_3en = buff_3en + 1
buff(buff_3en,buff_id) = w( il ,
2,i)
enddo
enddo

buff(i1:buff_3en,buff_id)=1987000000.d0,d0,d0
buff(i1:buff_3en,buff_id)

else if( dir .eq. -1 ) then

do i=0,n2-1
do il=i,n1
buff_3en = buff_3en + 1
buff(buff_3en,buff_id) = buff_id + 31.0d0-
1.0d1
enddo
enddo

buff(i1:buff_3en,buff_id)=1987000000.d0,d0,d0
buff(i1:buff_3en,buff_id)

endif
endif

if( axis .eq. 3 ) then
if( dir .eq. -2 ) then

do i=0,n2
do il=i,n1
buff_3en = buff_3en + 1
buff(buff_3en,buff_id) = w(
il,i2,n3-1)
enddo
enddo

buff(i1:buff_3en,buff_id)=1987000000.d0,d0,d0
buff(i1:buff_3en,buff_id)

else if( dir .eq. -1 ) then

do i=0,n2
do il=i,n1
buff_3en = buff_3en + 1
buff(buff_3en,buff_id) = w(
il,i2,n3-1)
enddo
enddo

buff(i1:buff_3en,buff_id)=1987000000.d0,d0,d0
buff(i1:buff_3en,buff_id)

endif
endif

return
end

subroutine tabul( axis, dir, n, n1, n2, n3 )
use def_variables

implicit none

include "config.h"
include "global.h"

integer axis, dir, N1, N2, N3
double precision w( N1, N2, N3 )

integer buff_id, iode

integer i3, i2, i1

buff_id = 3 + dir
iode = 0

if( axis .eq. 3 ) then
if( dir .eq. -2 ) then

do i=0,n2-1
do i2=n2,n1+1
w = iode + 1

```

```

    a[n3, i2, i3] = buff[i2size, buff_i4]
    memdb()

    if( axis[i] dir == 0) then
      do 1300, n3-1
        do 1200, n2-1
          a[n2, i2] = a[n2, i2+1]
          a[n3, i2, i3] = buff[i2size, buff_i4]
          memdb()
        enddo
      enddo
    else if( axis == 2) then
      if( dir == -1) then
        do 1300, n3-1
          do 1200, n2
            a[n2, i2] = a[n2, i2+1]
            a[n3, i2, i3] = buff[i2size, buff_i4]
            memdb()
          enddo
        enddo
      else if( dir == +1) then
        do 1300, n3-1
          do 1200, n2
            a[n2, i2] = a[n2, i2+1]
            a[n3, i2, i3] = buff[i2size, buff_i4]
            memdb()
          enddo
        enddo
      enddo
    else if( axis == 3) then
      if( dir == -1) then
        do 1400, n3
          do 1300, n2
            a[n2, i2] = a[n2, i2+1]
            a[n3, i2, i3] = buff[i2size, buff_i4]
            memdb()
          enddo
        enddo
      else if( dir == +1) then
        do 1400, n3
          do 1300, n2
            a[n2, i2] = a[n2, i2+1]
            a[n3, i2, i3] = buff[i2size, buff_i4]
            memdb()
          enddo
        enddo
      enddo
    enddo

    return
end

subroutine main[a] n1, n2, n3, n4
  use mdf_intrinsic

implicit none

include 'cstypes.h'
include 'global.h'

integer a(n1, n2, n3)
character precision(81), n1, n2, n3

integer i3, i2, i1, buff_i2size, buff_i4
integer i, kb, a(n1)

dir = -1

buff_i2 = 3 + dir
buff_i4 = kb2

```

[illegible]

```

buff_id = 3 + dir
lenm = 0

if( axis_eq, 2 ) then
  do i=1,n2-1
    lenm = lenm + 1
    a(i,12,3) = buff(i)lenm, buff_id
  enddo
endif

endif

if( axis_eq, 2 ) then
  do i=1,n2-1
    lenm = lenm + 1
    a(i,12,3) = buff(i)lenm, buff_id
  enddo
endif

+001.7

if( axis_eq, 3 ) then
  do i=1,n1
    lenm = lenm + 3
    a(i,12,3) = buff(i)lenm, buff_id
  enddo
endif

+001.7

dir = 4

buff_id = 3 + dir
lenm = 0

if( axis_eq, 3 ) then
  do i=1,n2-1
    lenm = lenm + 1
    a(i,12,3) = buff(i)lenm, buff_id
  enddo
endif

+001.7

if( axis_eq, 2 ) then
  do i=1,n1
    lenm = lenm + 3
    a(i,12,3) = buff(i)lenm, buff_id
  enddo
endif

+001.7

if( axis_eq, 3 ) then
  do i=1,n1
    lenm = lenm + 3
    a(i,12,3) = buff(i)lenm, buff_id
  enddo
endif

return

```



MPI Reality

Tramonto
WJDC
Functional

```

1  def __init__(self, num_embeddings: int, embedding_dim: int, padding_idx: Optional[int] = None):
2      super().__init__()
3      self.num_embeddings = num_embeddings
4      self.embedding_dim = embedding_dim
5      self.padding_idx = padding_idx
6      self._weight = None
7      self._padding_idx = None
8      self._num_embeddings = None
9      self._embedding_dim = None
10     self._padding_idx = None
11
12     def _check_input(self, input: torch.Tensor):
13         if not isinstance(input, torch.Tensor):
14             raise ValueError("Input must be a torch.Tensor")
15         if input.dim() != 2:
16             raise ValueError("Input must be a 2D tensor")
17         if input.size(1) != self.embedding_dim:
18             raise ValueError("Input must have embedding dimension {}".format(self.embedding_dim))
19
20     def forward(self, input: torch.Tensor) -> torch.Tensor:
21         self._check_input(input)
22         return torch.nn.functional.embedding(input, self._weight, self._padding_idx)
23
24     @property
25     def weight(self) -> torch.Tensor:
26         return self._weight
27
28     @weight.setter
29     def weight(self, weight: torch.Tensor) -> None:
30         self._weight = weight
31
32     @property
33     def padding_idx(self) -> Optional[int]:
34         return self._padding_idx
35
36     @padding_idx.setter
37     def padding_idx(self, padding_idx: Optional[int]) -> None:
38         self._padding_idx = padding_idx
39
40     @property
41     def num_embeddings(self) -> int:
42         return self._num_embeddings
43
44     @num_embeddings.setter
45     def num_embeddings(self, num_embeddings: int) -> None:
46         self._num_embeddings = num_embeddings
47
48     @property
49     def embedding_dim(self) -> int:
50         return self._embedding_dim
51
52     @embedding_dim.setter
53     def embedding_dim(self, embedding_dim: int) -> None:
54         self._embedding_dim = embedding_dim
55
56     def __repr__(self) -> str:
57         return "Embedding({} x {})".format(self.num_embeddings, self.embedding_dim)
58
59     def __str__(self) -> str:
60         return "Embedding({} x {})".format(self.num_embeddings, self.embedding_dim)
61
62     def __len__(self) -> int:
63         return self.num_embeddings
64
65     def __getitem__(self, index: int) -> torch.Tensor:
66         return self._weight[index]
67
68     def __setitem__(self, index: int, value: torch.Tensor) -> None:
69         self._weight[index] = value
70
71     def __delitem__(self, index: int) -> None:
72         del self._weight[index]
73
74     def __iter__(self) -> torch.Tensor:
75         return iter(self._weight)
76
77     def __contains__(self, index: int) -> bool:
78         return 0 <= index < self.num_embeddings
79
80     def __add__(self, other: torch.Tensor) -> torch.Tensor:
81         return self._weight + other
82
83     def __sub__(self, other: torch.Tensor) -> torch.Tensor:
84         return self._weight - other
85
86     def __mul__(self, other: torch.Tensor) -> torch.Tensor:
87         return self._weight * other
88
89     def __div__(self, other: torch.Tensor) -> torch.Tensor:
90         return self._weight / other
91
92     def __mod__(self, other: torch.Tensor) -> torch.Tensor:
93         return self._weight % other
94
95     def __pow__(self, other: torch.Tensor) -> torch.Tensor:
96         return self._weight ** other
97
98     def __matmul__(self, other: torch.Tensor) -> torch.Tensor:
99         return self._weight @ other
100    def __rmatmul__(self, other: torch.Tensor) -> torch.Tensor:
101        return other @ self._weight
102
103    def __truediv__(self, other: torch.Tensor) -> torch.Tensor:
104        return self._weight / other
105
106    def __rtruediv__(self, other: torch.Tensor) -> torch.Tensor:
107        return other / self._weight
108
109    def __divmod__(self, other: torch.Tensor) -> (torch.Tensor, torch.Tensor):
110        return self._weight // other, self._weight % other
111
112    def __rdivmod__(self, other: torch.Tensor) -> (torch.Tensor, torch.Tensor):
113        return other // self._weight, other % self._weight
114
115    def __neg__(self) -> torch.Tensor:
116        return -self._weight
117
118    def __pos__(self) -> torch.Tensor:
119        return self._weight
120
121    def __abs__(self) -> torch.Tensor:
122        return self._weight.abs()
123
124    def __eq__(self, other: torch.Tensor) -> bool:
125        return self._weight == other
126
127    def __neq__(self, other: torch.Tensor) -> bool:
128        return self._weight != other
129
130    def __lt__(self, other: torch.Tensor) -> bool:
131        return self._weight < other
132
133    def __le__(self, other: torch.Tensor) -> bool:
134        return self._weight <= other
135
136    def __gt__(self, other: torch.Tensor) -> bool:
137        return self._weight > other
138
139    def __ge__(self, other: torch.Tensor) -> bool:
140        return self._weight >= other
141
142    def __and__(self, other: torch.Tensor) -> torch.Tensor:
143        return self._weight & other
144
145    def __or__(self, other: torch.Tensor) -> torch.Tensor:
146        return self._weight | other
147
148    def __xor__(self, other: torch.Tensor) -> torch.Tensor:
149        return self._weight ^ other
150
151    def __lshift__(self, other: torch.Tensor) -> torch.Tensor:
152        return self._weight << other
153
154    def __rshift__(self, other: torch.Tensor) -> torch.Tensor:
155        return self._weight >> other
156
157    def __invert__(self) -> torch.Tensor:
158        return ~self._weight
159
160    def __bitwise_and__(self, other: torch.Tensor) -> torch.Tensor:
161        return self._weight & other
162
163    def __bitwise_or__(self, other: torch.Tensor) -> torch.Tensor:
164        return self._weight | other
165
166    def __bitwise_xor__(self, other: torch.Tensor) -> torch.Tensor:
167        return self._weight ^ other
168
169    def __lshift__(self, other: torch.Tensor) -> torch.Tensor:
170        return self._weight << other
171
172    def __rshift__(self, other: torch.Tensor) -> torch.Tensor:
173        return self._weight >> other
174
175    def __invert__(self) -> torch.Tensor:
176        return ~self._weight
177
178    def __bitwise_and__(self, other: torch.Tensor) -> torch.Tensor:
179        return self._weight & other
180
181    def __bitwise_or__(self, other: torch.Tensor) -> torch.Tensor:
182        return self._weight | other
183
184    def __bitwise_xor__(self, other: torch.Tensor) -> torch.Tensor:
185        return self._weight ^ other
186
187    def __lshift__(self, other: torch.Tensor) -> torch.Tensor:
188        return self._weight << other
189
190    def __rshift__(self, other: torch.Tensor) -> torch.Tensor:
191        return self._weight >> other
192
193    def __invert__(self) -> torch.Tensor:
194        return ~self._weight
195
196    def __bitwise_and__(self, other: torch.Tensor) -> torch.Tensor:
197        return self._weight & other
198
199    def __bitwise_or__(self, other: torch.Tensor) -> torch.Tensor:
200        return self._weight | other
201
202    def __bitwise_xor__(self, other: torch.Tensor) -> torch.Tensor:
203        return self._weight ^ other
204
205    def __lshift__(self, other: torch.Tensor) -> torch.Tensor:
206        return self._weight << other
207
208    def __rshift__(self, other: torch.Tensor) -> torch.Tensor:
209        return self._weight >> other
210
211    def __invert__(self) -> torch.Tensor:
212        return ~self._weight
213
214    def __bitwise_and__(self, other: torch.Tensor) -> torch.Tensor:
215        return self._weight & other
216
217    def __bitwise_or__(self, other: torch.Tensor) -> torch.Tensor:
218        return self._weight | other
219
220    def __bitwise_xor__(self, other: torch.Tensor) -> torch.Tensor:
221        return self._weight ^ other
222
223    def __lshift__(self, other: torch.Tensor) -> torch.Tensor:
224        return self._weight << other
225
226    def __rshift__(self, other: torch.Tensor) -> torch.Tensor:
227        return self._weight >> other
228
229    def __invert__(self) -> torch.Tensor:
230        return ~self._weight
231
232    def __bitwise_and__(self, other: torch.Tensor) -> torch.Tensor:
233        return self._weight & other
234
235    def __bitwise_or__(self, other: torch.Tensor) -> torch.Tensor:
236        return self._weight | other
237
238    def __bitwise_xor__(self, other: torch.Tensor) -> torch.Tensor:
239        return self._weight ^ other
240
241    def __lshift__(self, other: torch.Tensor) -> torch.Tensor:
242        return self._weight << other
243
244    def __rshift__(self, other: torch.Tensor) -> torch.Tensor:
245        return self._weight >> other
246
247    def __invert__(self) -> torch.Tensor:
248        return ~self._weight
249
250    def __bitwise_and__(self, other: torch.Tensor) -> torch.Tensor:
251        return self._weight & other
252
253    def __bitwise_or__(self, other: torch.Tensor) -> torch.Tensor:
254        return self._weight | other
255
256    def __bitwise_xor__(self, other: torch.Tensor) -> torch.Tensor:
257        return self._weight ^ other
258
259    def __lshift__(self, other: torch.Tensor) -> torch.Tensor:
260        return self._weight << other
261
262    def __rshift__(self, other: torch.Tensor) -> torch.Tensor:
263        return self._weight >> other
264
265    def __invert__(self) -> torch.Tensor:
266        return ~self._weight
267
268    def __bitwise_and__(self, other: torch.Tensor) -> torch.Tensor:
269        return self._weight & other
270
271    def __bitwise_or__(self, other: torch.Tensor) -> torch.Tensor:
272        return self._weight | other
273
274    def __bitwise_xor__(self, other: torch.Tensor) -> torch.Tensor:
275        return self._weight ^ other
276
277    def __lshift__(self, other: torch.Tensor) -> torch.Tensor:
278        return self._weight << other
279
280    def __rshift__(self, other: torch.Tensor) -> torch.Tensor:
281        return self._weight >> other
282
283    def __invert__(self) -> torch.Tensor:
284        return ~self._weight
285
286    def __bitwise_and__(self, other: torch.Tensor) -> torch.Tensor:
287        return self._weight & other
288
289    def __bitwise_or__(self, other: torch.Tensor) -> torch.Tensor:
290        return self._weight | other
291
292    def __bitwise_xor__(self, other: torch.Tensor) -> torch.Tensor:
293        return self._weight ^ other
294
295    def __lshift__(self, other: torch.Tensor) -> torch.Tensor:
296        return self._weight << other
297
298    def __rshift__(self, other: torch.Tensor) -> torch.Tensor:
299        return self._weight >> other
300
301    def __invert__(self) -> torch.Tensor:
302        return ~self._weight
303
304    def __bitwise_and__(self, other: torch.Tensor) -> torch.Tensor:
305        return self._weight & other
306
307    def __bitwise_or__(self, other: torch.Tensor) -> torch.Tensor:
308        return self._weight | other
309
310    def __bitwise_xor__(self, other: torch.Tensor) -> torch.Tensor:
311        return self._weight ^ other
312
313    def __lshift__(self, other: torch.Tensor) -> torch.Tensor:
314        return self._weight << other
315
316    def __rshift__(self, other: torch.Tensor) -> torch.Tensor:
317        return self._weight >> other
318
319    def __invert__(self) -> torch.Tensor:
320        return ~self._weight
321
322    def __bitwise_and__(self, other: torch.Tensor) -> torch.Tensor:
323        return self._weight & other
324
325    def __bitwise_or__(self,
```

[illegible]

- New functional.
- Bonded systems.
- 552 lines C code.

How much MPI-specific code?

dft_fill_wjdc.c
MPI-specific
code

[illegible][illegible]

```

A,MIH,P,C,LOCAL,2)=ZERO
A,MIH,P,C,LOCAL,MI)=ZERO
A,MIH,P,C,LOCAL,N1)=ZERO
A,MIH,P,C,LOCAL,N2)=ZERO
A,MIH,P,C,LOCAL,2)=ZERO
A,MIH,P,C,LOCAL,MI)=ZERO
A,MIH,P,C,LOCAL,N1)=ZERO
A,MIH,P,C,LOCAL,N2)=ONE
E,MIH,P,C,LOCAL,2)=ONE
ENDDOF
ENDIF
1
CFEM & STAT,best (nan wit)
IF LOCAL,0)=CALL,STAT77(1)
SIM,ALL=SIM,S,C,comp
SIM,ALL=SIM,S,C,comp
ENDIF
CFEM & STAT,end (nan wit)
call unset_all_error

RETURN
END SUBROUTINE SOURCE_FF,C

// Comments on the modifications for DMP version implementation
// 001 Includes header file and common declarations for parallelization
// 002 Includes de long lines 1,0,found=2 (0,0,0),0,found1

```

- MPI-callable, OpenMP-enabled.
- 340 Fortran lines.
- No MPI-specific code.
- Ubiquitous OpenMP markup (red regions).

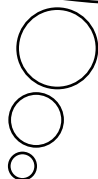
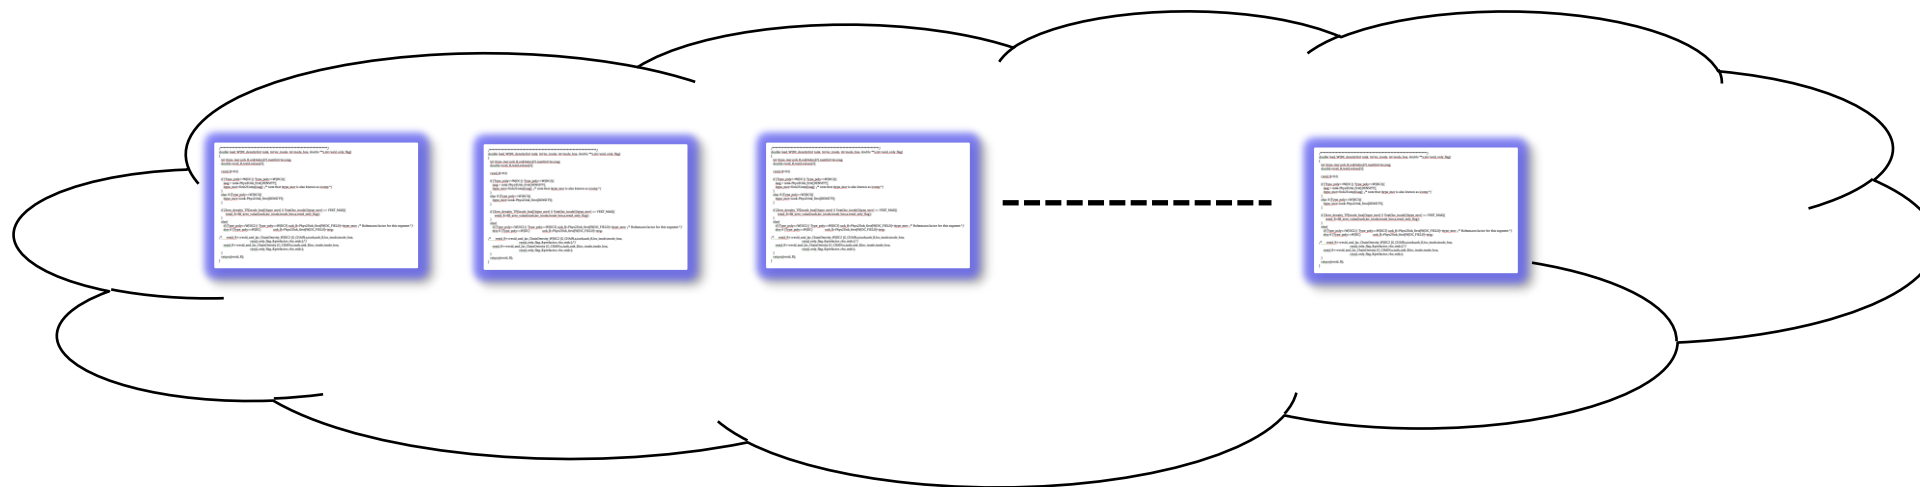
MFIX: Multiphase Flows with Interphase eXchanges (<https://www.mfix.org/>)



Reasons for MPI Success?

- Portability? Yes.
- Standardized? Yes.
- Momentum? Yes.
- Separation of many
Parallel & Algorithms
concerns? Big Yes.
- Once framework in place:
 - Sophisticated physics added as serial code.
 - Ratio of science experts vs. parallel experts: 10:1.
- Key goal for new parallel apps: Preserve this ratio

Computational Domain Expert Writing MPI Code



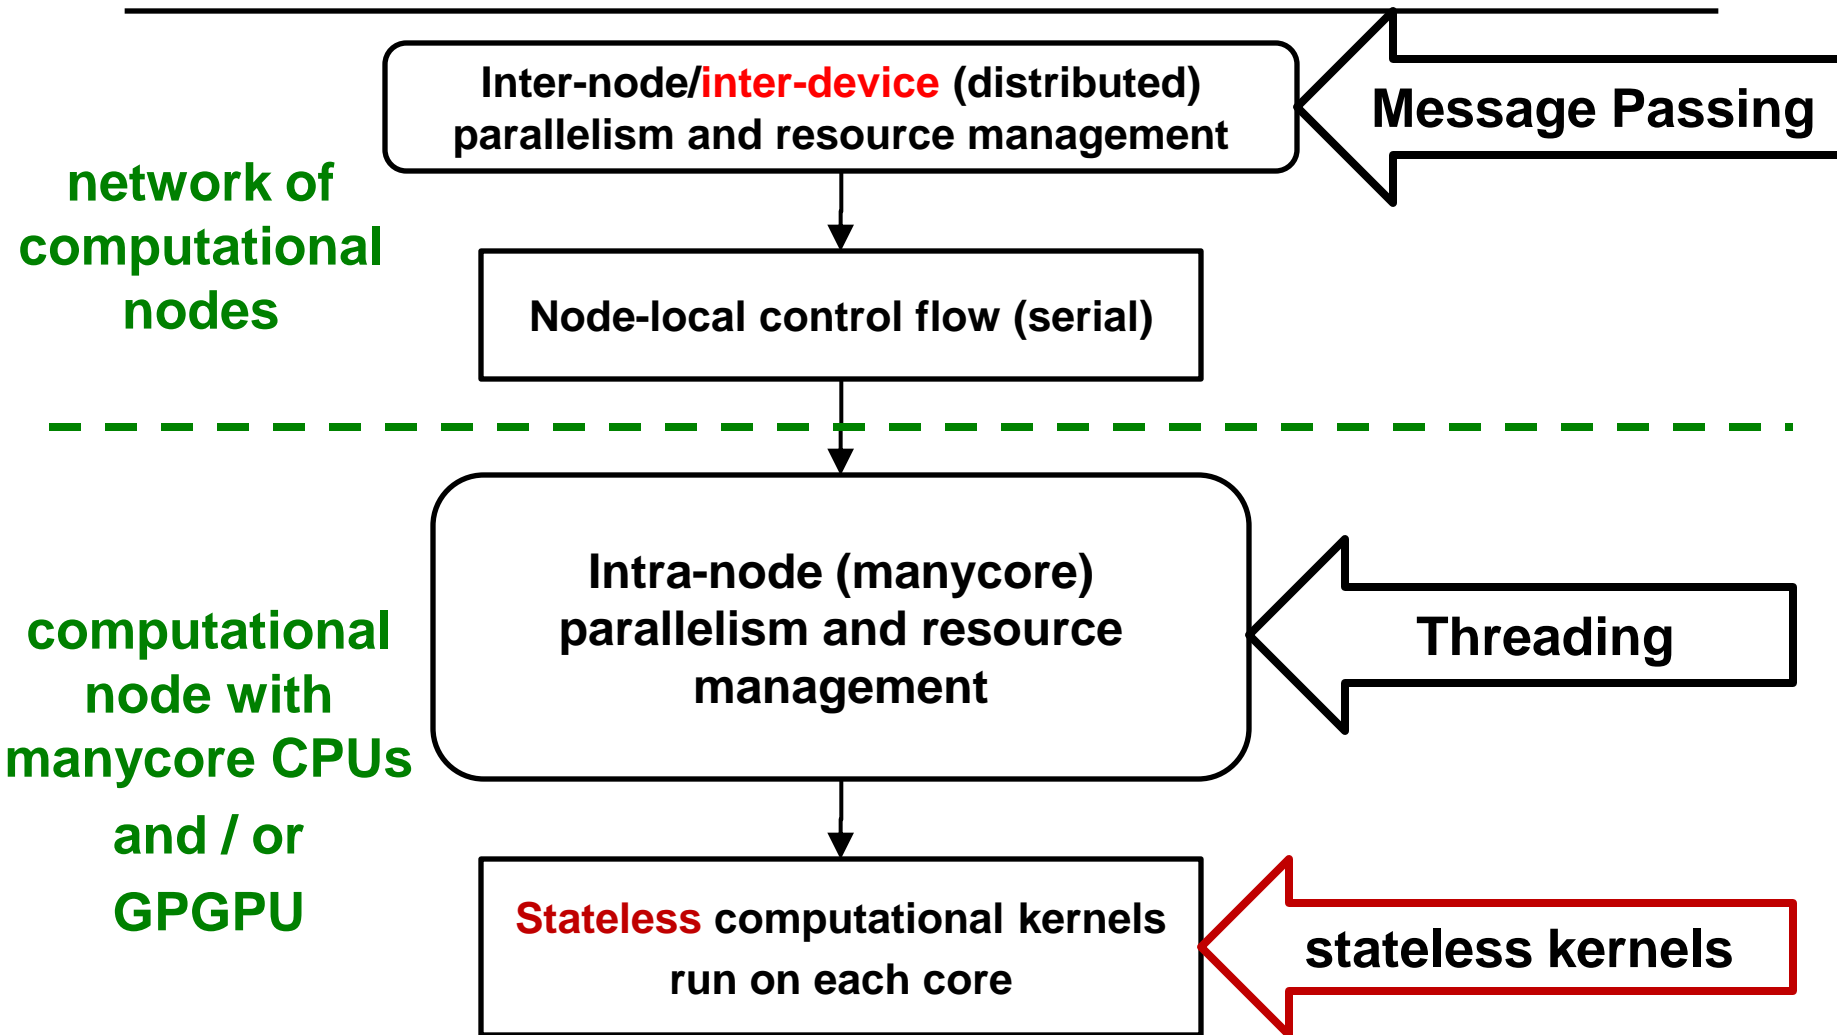
Computational Domain Expert Writing Future Parallel Code





Evolving Parallel Programming Model

Parallel Programming Model: Multi-level/Multi-device



Adapted from slide of H. Carter Edwards



Domain Scientist's Parallel Palette

- MPI-only (SPMD) apps:
 - Single parallel construct.
 - Simultaneous execution.
 - Parallelism of even the messiest serial code.
- Next-generation applications:
 - Internode:
 - MPI, yes, or something like it.
 - Composed with intranode.
 - Intranode:
 - Much richer palette.
 - More care required from programmer.
- What are the constructs in our new palette?



Obvious Constructs/Concerns

- Parallel for:
 - No loop-carried dependence.
 - Rich loops.
- Parallel reduce:
 - Couple with other computations.
 - Concern for reproducibility.




Other construct: Pipeline

- Sequence of filters.
- Each filter is:
 - Sequential (grab element ID, enter global assembly) or
 - Parallel (fill element stiffness matrix).
- Filters executed in sequence.
- Programmer's concern:
 - Determine (conceptually): Can filter execute in parallel?
 - Write filter (serial code).
 - Register it with the pipeline.
- Extensible:
 - New physics feature.
 - New filter added to pipeline.



Other construct: Thread team

- Multiple threads.
- Fast barrier.
- Shared, fast access memory pool.
- Example: Nvidia SM
- X86 more vague, emerging more clearly in future.

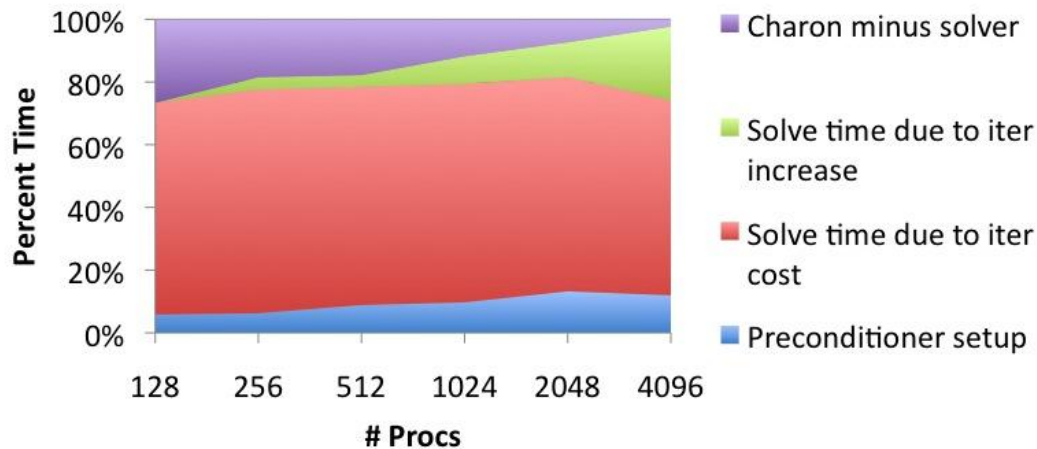


Finite Elements/Volumes/Differences and parallel node constructs

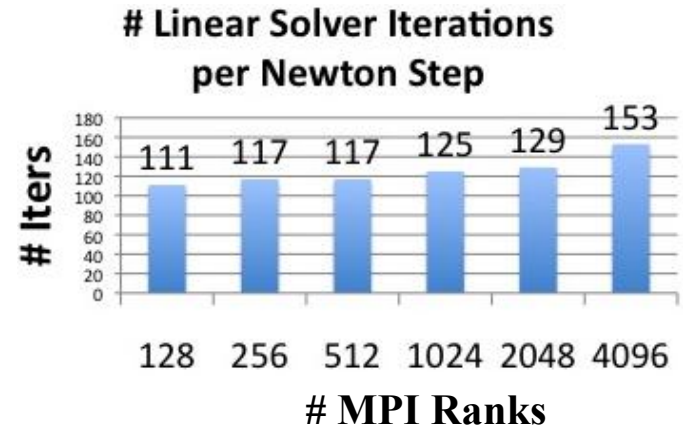
- Parallel for, reduce, pipeline:
 - Sufficient for vast majority of node level computation.
 - Supports:
 - Complex modeling expression.
 - Vanilla parallelism.
- Thread team:
 - Complicated.
 - Requires true parallel algorithm knowledge.
 - Useful in solvers.

Preconditioners for Scalable Multicore Systems

Charon Timing Breakdown on TLCC
Strong Scaling 28M Unknowns



Strong scaling of Charon on TLCC (P. Lin, J. Shadid 2009)

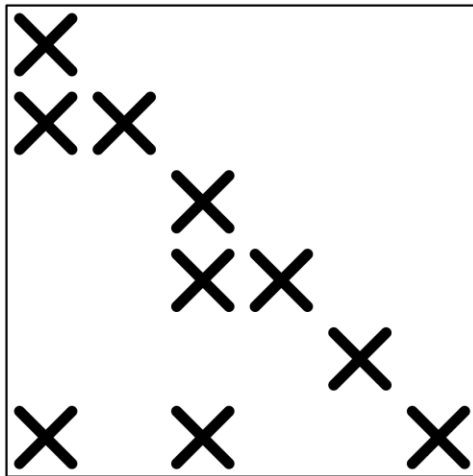


- Observe: Iteration count increases with number of subdomains.
- With scalable threaded triangular solves
 - Solve triangular system on larger subdomains.
 - Reduce number of subdomains.
- Goal:
 - Better kernel scaling (threads vs. MPI processes).
 - Better convergence, More robust.
- Note: App (-solver) scales very well in MPI-only mode.
- Exascale Potential: Tiled, pipelined implementation.

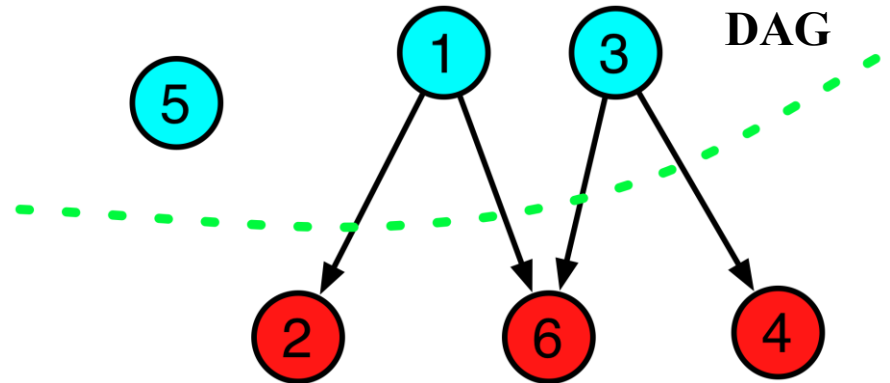
MPI Tasks	Threads	Iterations
4096	1	153
2048	2	129
1024	4	125
512	8	117
256	16	117
128	32	111

Factors Impacting Performance of Multithreaded Sparse Triangular Solve, Michael M. Wolf and Michael A. Heroux and Erik G. Boman, VECPAR 2010, to appear.

Level Set Triangular Solver



L



Triangular Solve:

- Critical Kernel
 - MG Smoothers
 - Incomplete IC/ILU
- Naturally Sequential
- Building on classic algorithms:
 - Level Sched:
 - circa 1990.
 - Vectorization.
 - New: Generalized.

$$\tilde{L} = P L P^T = \begin{bmatrix} D_1 & & & & \\ A_{2,1} & D_2 & & & \\ A_{3,1} & A_{3,2} & D_3 & & \\ \vdots & \vdots & \vdots & \ddots & \\ A_{l,1} & A_{l,2} & A_{l,3} & \dots & D_l \end{bmatrix} \quad \text{Permuted System}$$

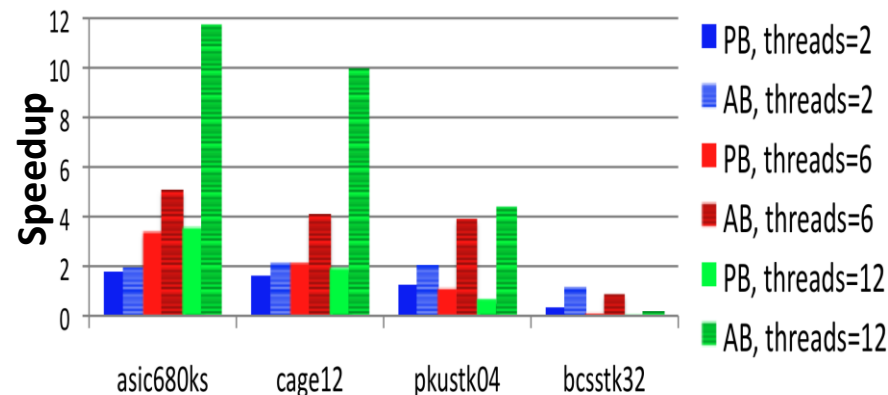
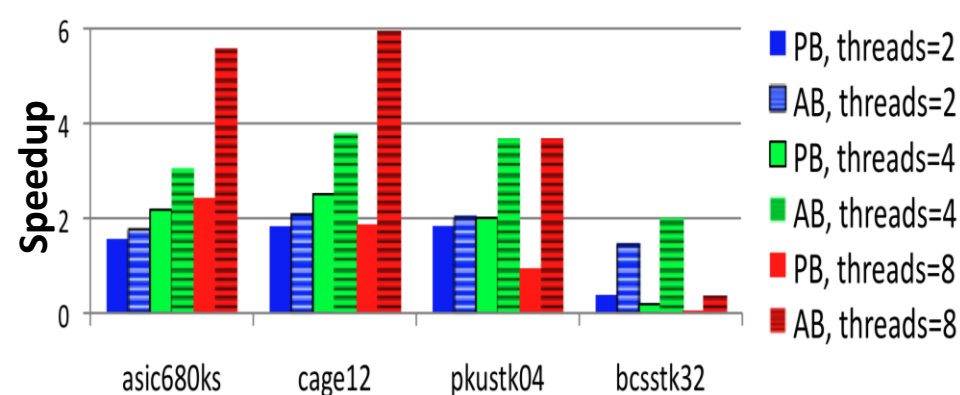
$$\begin{aligned} \tilde{x}_1 &= D_1^{-1} \tilde{y}_1 \\ \tilde{x}_2 &= D_2^{-1} (\tilde{y}_2 - A_{2,1} \tilde{x}_1) \\ \vdots &\quad \quad \quad \vdots \\ \tilde{x}_l &= D_l^{-1} (\tilde{y}_l - A_{l,1} \tilde{x}_1 - \dots - A_{l,l-1} \tilde{x}_{l-1}) \end{aligned} \quad \begin{array}{c} \text{Multi-step} \\ \text{Algorithm} \end{array}$$



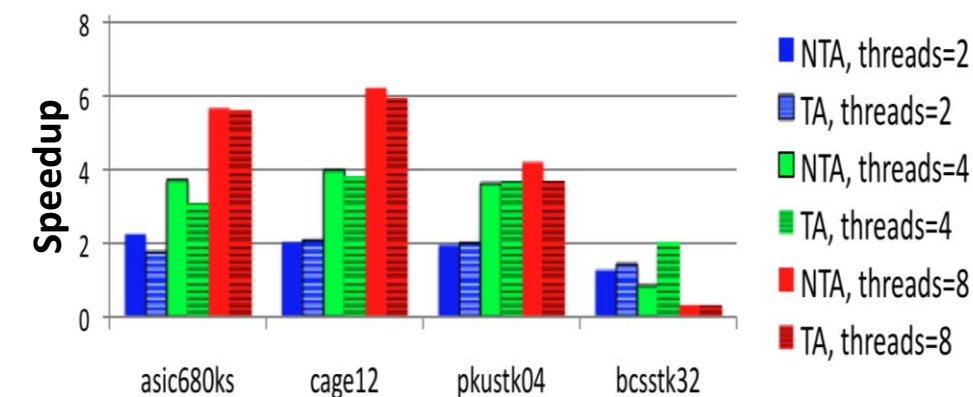
Triangular Solve Results

Name	N	nnz	N/nlevels	application area
asic680ks	682,712	2,329,176	13932.9	circuit simulation
cage12	130,228	2,032,536	1973.2	DNA electrophoresis
pkustk04	55,590	4,218,660	149.4	structural engineering
bcsstk32	44,609	2,014,701	15.1	structural engineering

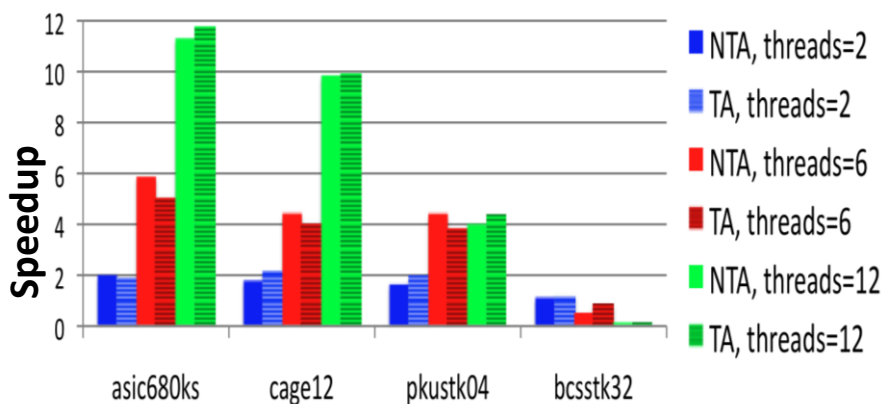
Passive (PB) vs. Active (AB) Barriers: Critical for Performance



Nehalem



Istanbul



AB + No Thread Affinity (NTA) vs. AB + Thread Affinity (TA) : Also Helpful



Thread Team Advantages

- **Qualitatively better algorithm:**
 - Threaded triangular solve scales.
 - Fewer MPI ranks means fewer iterations, better robustness.
- **Exploits:**
 - Shared data.
 - Fast barrier.
 - Data-driven parallelism.



Placement and Migration



Placement and Migration

- MPI:
 - Data/work placement clear.
 - Migration explicit.
- Threading:
 - It's a mess (IMHO).
 - Some platforms good.
 - Many not.
 - Default is bad (but getting better).
 - Some issues are intrinsic.



Data Placement on NUMA

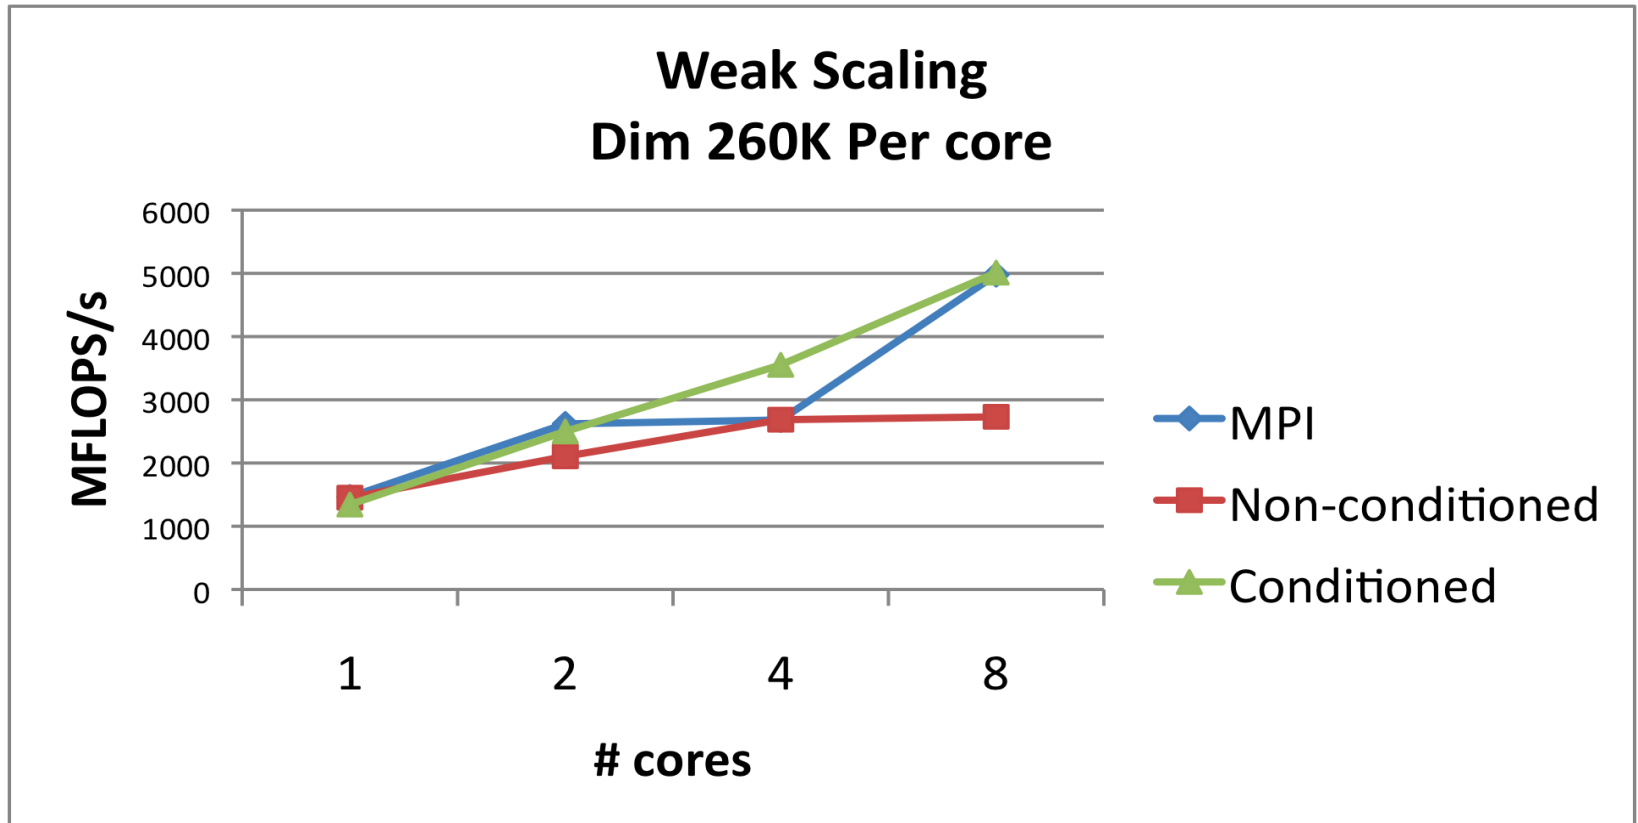
- Memory Intensive computations: Page placement has huge impact.
- Most systems: First touch (except LWKs).
- Application data objects:
 - Phase 1: Construction phase, e.g., finite element assembly.
 - Phase 2: Use phase, e.g., linear solve.
- Problem: First touch difficult to control in phase 1.
- Idea: Page migration.
 - Not new: SGI Origin. Many old papers on topic.



Data placement experiments

- MiniApp: HPCCG (Mantevo Project)
- Construct sparse linear system, solve with CG.
- Two modes:
 - Data placed by assembly, not migrated for NUMA
 - Data migrated using parallel access pattern of CG.
- Results on dual socket quad-core Nehalem system.

Weak Scaling Problem



- MPI and conditioned data approach comparable.
- Non-conditioned very poor scaling.



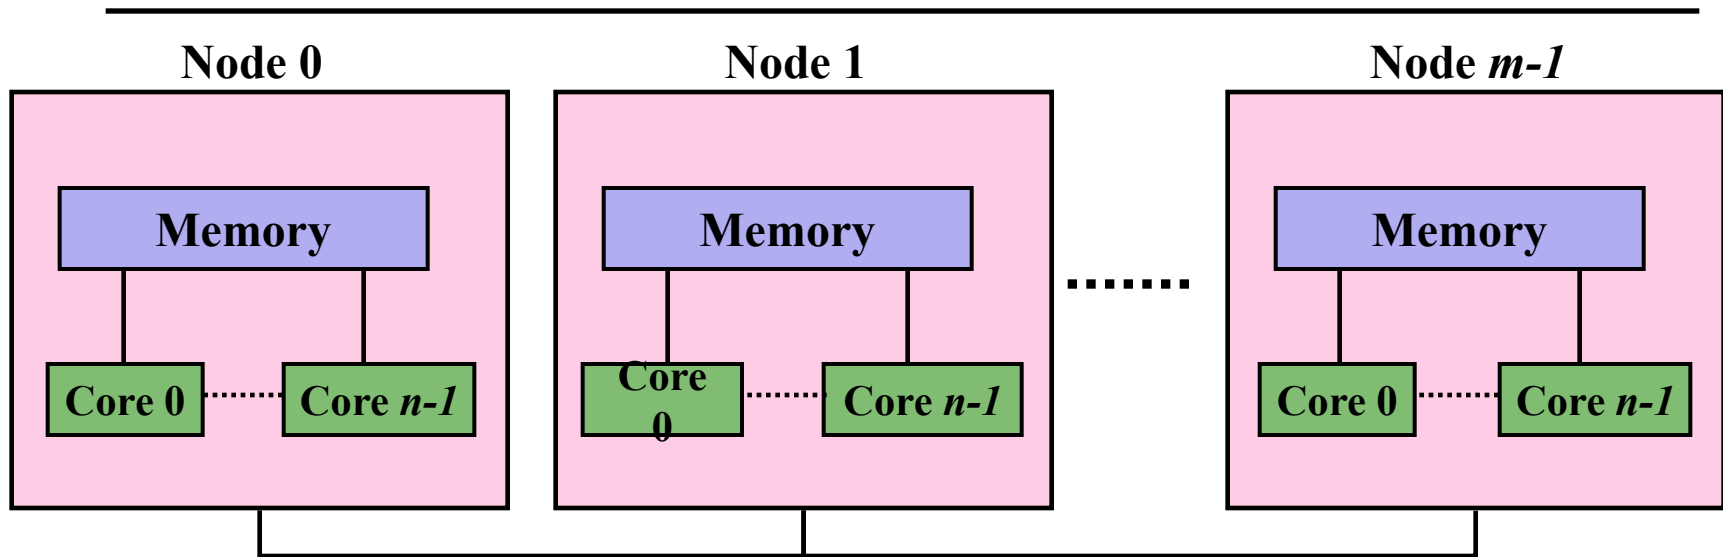
Page Placement summary

- MPI+OpenMP (or any threading approach) is best overall.
- But:
 - Data placement is big issue.
 - Hard to control.
 - Insufficient runtime support.
- Current work:
 - Migrate on next-touch (MONT).
 - Considered in OpenMP (next version).
 - Also being studied in Kitten (Kevin Pedretti).
- Note: This phenomenon especially damaging to OpenMP common usage.



Transition: MPI-only to MPI+[X/Y/Z]

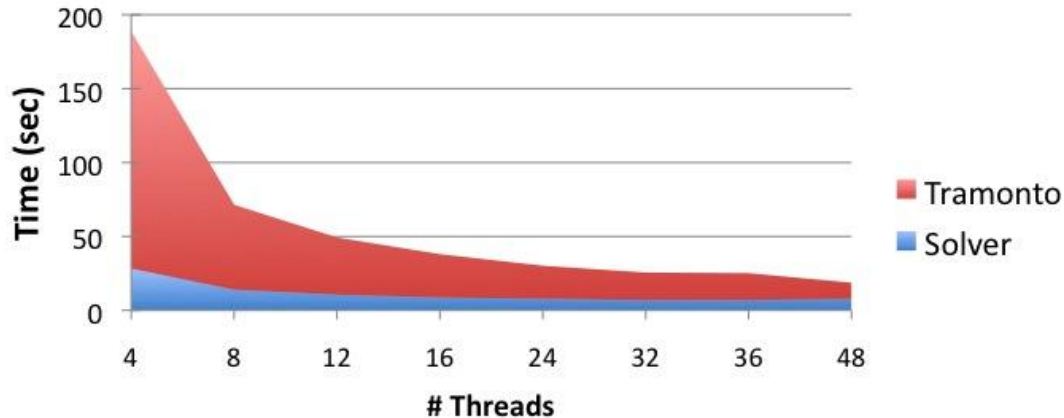
Parallel Machine Block Diagram



- Parallel machine with $p = m * n$ processors:
 - m = number of nodes.
 - n = number of shared memory processors per node.
- Two ways to program:
 - Way 1: p MPI processes.
 - Way 2: m MPI processes with n threads per MPI process.
- New third way:
 - “Way 1” in some parts of the execution (the app).
 - “Way 2” in others (the solver).

Multicore Scaling: App vs. Solver

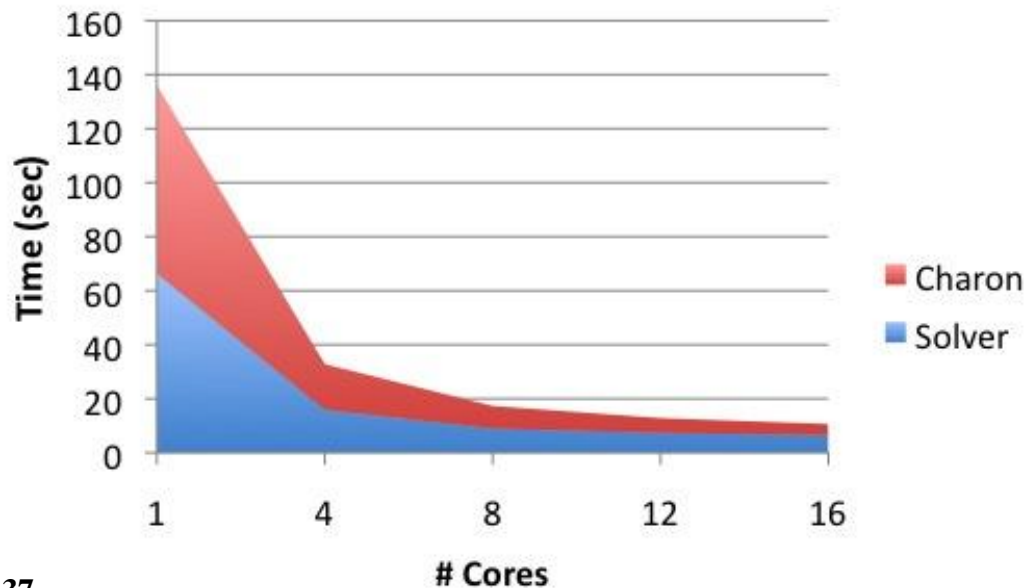
Tramonto vs. Solver Time on Niagara2:
4-48 Threads



Application:

- Scales well (sometimes superlinear)
- MPI-only sufficient.

Charon vs Solver Time: 1-16 Cores

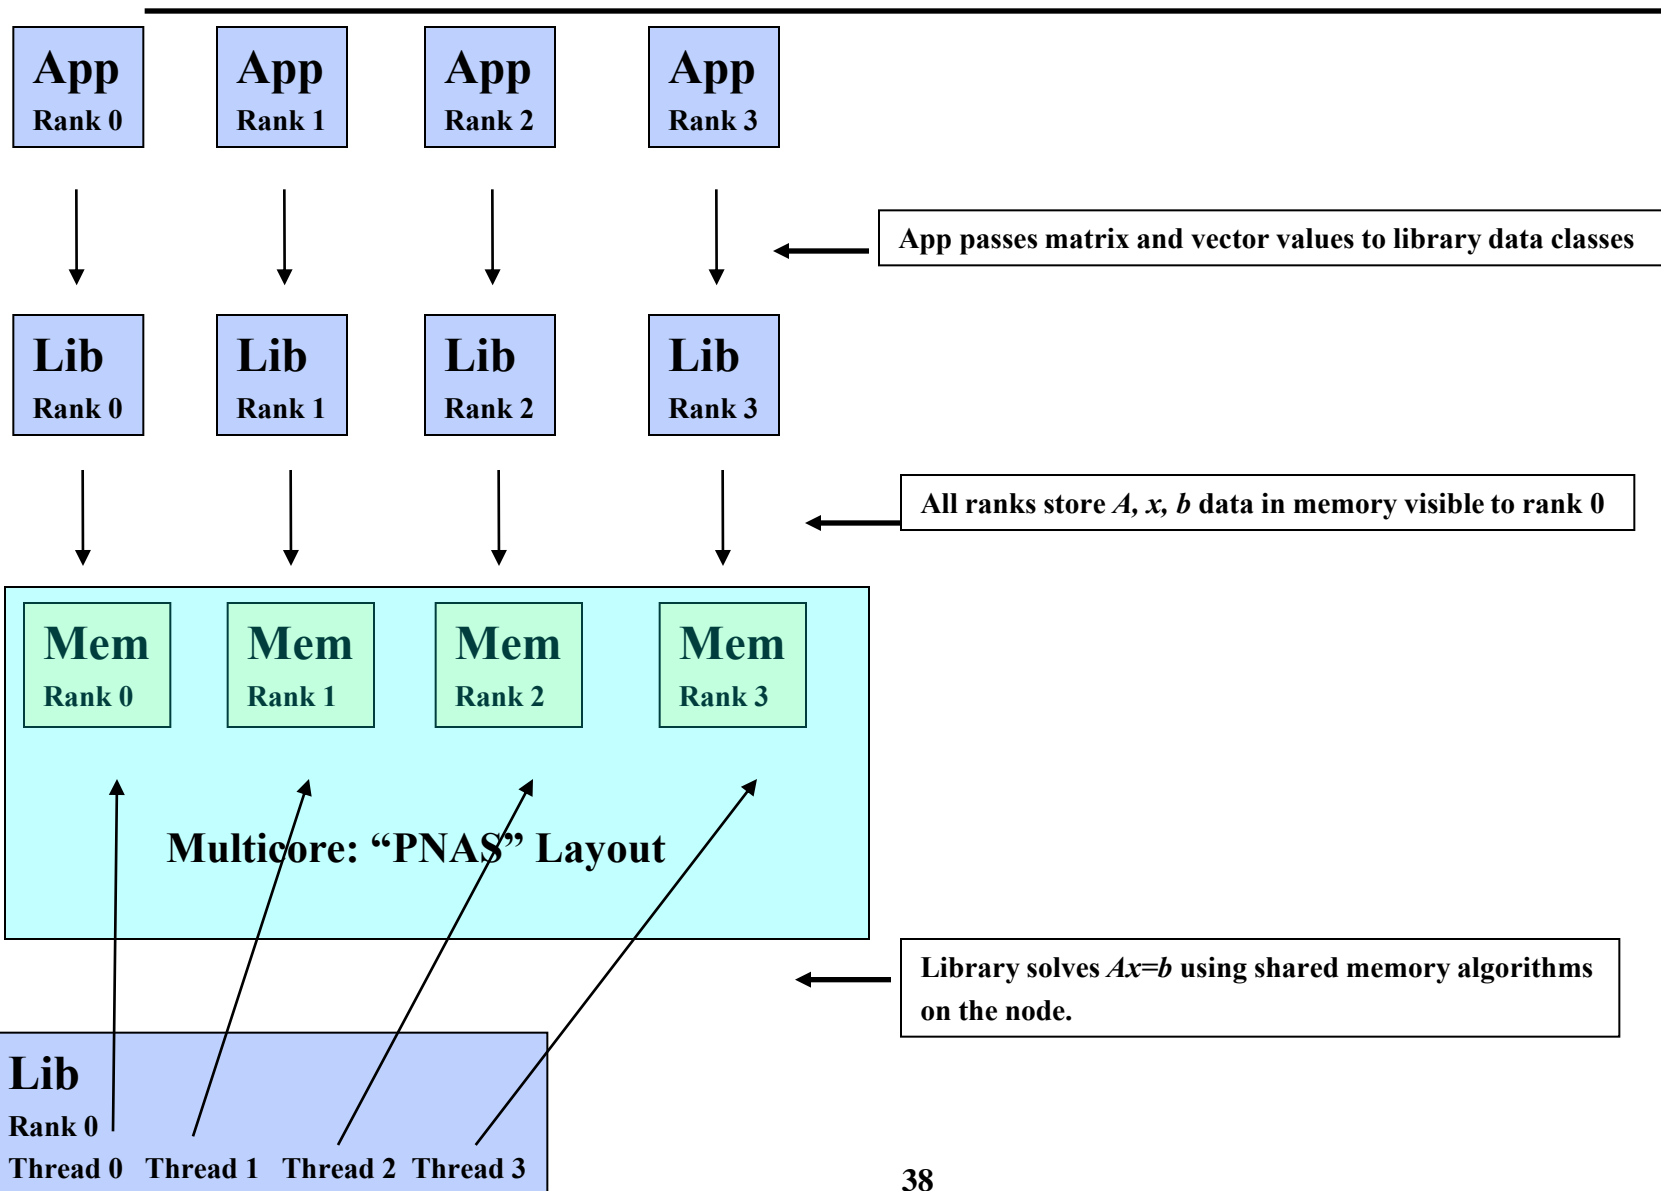


Solver:

- Scales more poorly.
- Memory system-limited.
- MPI+threads can help.

* Charon Results:
Lin & Shadid TLCC Report

MPI-Only + MPI/Threading: $Ax=b$



MPI Shared Memory Allocation

Idea:

- Shared memory alloc/free functions:
 - MPI_Comm_alloc_mem
 - MPI_Comm_free_mem
- Predefined communicators:
 - MPI_COMM_NODE – ranks on node
 - MPI_COMM_SOCKET – UMA ranks
 - MPI_COMM_NETWORK – inter node
- Status:
 - Available in current development branch of OpenMPI.
 - First “Hello World” Program works.
 - Incorporation into standard still not certain. Need to build case.
 - Next Step: Demonstrate usage with threaded triangular solve.
- Exascale potential:
 - Incremental path to MPI+X.
 - Dial-able SMP scope.

```
int n = ...;
double* values;
MPI_Comm_alloc_mem(
    MPI_COMM_NODE, // comm (SOCKET works too)
    n*sizeof(double), // size in bytes
    MPI_INFO_NULL, // placeholder for now
    &values);       // Pointer to shared array (out)


// At this point:
// - All ranks on a node/socket have pointer to a shared buffer (values).
// - Can continue in MPI mode (using shared memory algorithms) or
// - Can quiet all but one:
int rank;
MPI_Comm_rank(MPI_COMM_NODE, &rank);
if (rank==0) { // Start threaded code segment, only on rank 0 of the node
    ...
}

MPI_Comm_free_mem(MPI_COMM_NODE, values);
```

Collaborators: B. Barrett, Brightwell, Wolf - SNL; Vallee, Koenig - ORNL



Resilient Algorithms



My Luxury in Life (wrt FT/Resilience)

The privilege to think of a computer as a *reliable, digital* machine.

“At 8 nm process technology, it will be harder to tell a 1 from a 0.”

(W. Camp 2008, 2010)



Users' View of the System Now

- “All nodes up and running.”
- Certainly nodes fail, but invisible to user.
- No need for me to be concerned.
- Someone else’s problem.



Users' View of the System Future

- Nodes in one of four states.
 1. Dead.
 2. Dying (perhaps producing faulty results).
 3. Reviving.
 4. Running properly:
 - a) Fully reliable or...
 - b) Maybe still producing an occasional bad result.



Faults: Hard vs. Soft

- Hard:
 - Program flow interrupted.
 - Majority of faults.
 - Presently handled by (global) checkpoint/restart.
 - Numerous papers on alternatives.
- Soft:
 - Program flow continues.
 - Minor perturbations in data state:
 - Incorrect address lookup (but still in user scope).
 - Incorrect FP value.



Algorithm-Based (Hard) Fault Tolerance

- Numerous approaches.
- Most common strategies:
 - Meta data:
 - Embed meta data into user-defined data structures.
 - Manage fault detection, recovery manually.
 - Algorithm results validation:
 - Use known algorithm properties.
 - Validate computed to known (e.g., residual check).
- Note: A lack of app awareness.

Common Approach to FT (Diplomacy Analogy)

We have linearized our portion of the nonlinear problem and would like you to negotiate a global linear solution with the other processors.

Yes, Madame President. I will return with our portion of the global linear solution, ASAP.

Madame President, although there was some rough weather, our fault tolerant linear solver worked and I have returned with our portion of the linear solution.

Thank you, we recovered nonlinear state, the linear solution is expensive. We can use your results.





Hard Error Futures

- C/R will continue as dominant approach:
 - Global state to global file system OK for small systems.
 - Large systems: State control will be localized, use SSD.
- Checkpoint-less restart:
 - Requires full vertical HW/SW stack co-operation.
 - Very challenging.
 - Stratified research efforts not effective.



Soft Error Futures

- Soft error handling: A legitimate algorithms issue.
- Programming model, runtime environment play role.



Consider GMRES as an example of how soft errors affect correctness

- Basic Steps
 - 1) Compute Krylov subspace (preconditioned sparse matrix-vector multiplies)
 - 2) Compute orthonormal basis for Krylov subspace (matrix factorization)
 - 3) Compute vector yielding minimum residual in subspace (linear least squares)
 - 4) Map to next iterate in the full space
 - 5) Repeat until residual is sufficiently small
- More examples in Bronevetsky & Supinski, 2008



Why GMRES?

- Many apps are implicit.
- Most popular (nonsymmetric) linear solver is preconditioned GMRES.
- Only small subset of calculations need to be reliable.
 - GMRES is iterative, but also direct.



Every calculation matters

- Small PDE Problem: Dim 21K, Nz 923K.
- ILUT/GMRES
- Correct computation 35 lters: 343M FLOPS
- Two examples of a **single** bad floating point op

Description	Iterations	FLOPS	Recursive Residual Error	Solution Error
All Correct Calcs	35	343M	4.6e-15	1.0e-6
Iter=2, y[1] += 1.0 SpMV incorrect Ortho subspace	35	343M	6.7e-15	3.7e+3
Q[1][1] += 1.0 Non-ortho subspace	N/C	N/A	7.7e-02	5.9e+5



One possible approach is transactional computation

- Database transactions: atomic
- Transactional memory: atomic memory operation
- Transactional computation:
 - Designated sensitive computation region (orthogonalization step in GMRES)
 - Guarantee accurate computation or notify user.



Needs to be coupled with SW-enabled guaranteed data regions

- User-designated reliable data region
- Extra protection to improve reliable data storage and transfer
- Examples
 - Original input data (needed for verification)
 - Linear solver: A , x , b
 - Orthogonal vectors for GMRES
- OpenMP pragma-enabled?



Goal

- Algorithms well-conditioned wrt soft failure.
- Now:
 - Single soft error produces erroneous results.
- Goal:
 - Correct results always.
 - Cost increase proportional to number of soft errors.
- Note: These are just two approaches to ABFT.



Software Development and Delivery

Compile-time Polymorphism

Templates and Sanity upon a shifting foundation

Software delivery:

- Essential Activity

How can we:

- Implement mixed precision algorithms?
- Implement generic fine-grain parallelism?
- Support hybrid CPU/GPU computations?
- Support extended precision?
- Explore redundant computations?
- Prepare for both exascale “swim lanes”?

C++ templates only sane way:

- Moving to completely templated Trilinos libraries.
- Other important benefits.
- **A usable stack exists now in Trilinos.**

Template Benefits:

- Compile time polymorphism.
- True generic programming.
- No runtime performance hit.
- Strong typing for mixed precision.
- Support for extended precision.
- Many more...

Template Drawbacks:

- Huge compile-time performance hit:
 - But good use of multicore :)
 - Eliminated for common data types.
- Complex notation:
 - Esp. for Fortran & C programmers).
 - Can insulate to some extent.



Solver Software Stack



Phase I packages: SPMD, int/double

Phase II packages: Templated

Optimization Unconstrained: Constrained:	Find $u \in \mathbb{R}^n$ that minimizes $g(u)$ Find $x \in \mathbb{R}^m$ and $u \in \mathbb{R}^n$ that minimizes $g(x, u)$ s.t. $f(x, u) = 0$	Sensitivities (Automatic Differentiation: Sacado)	MOOCHO
Bifurcation Analysis	Given nonlinear operator $F(x, u) \in \mathbb{R}^{n+m}$ For $F(x, u) = 0$ find space $u \in U \ni \frac{\partial F}{\partial x}$		LOCA
Transient Problems DAEs/ODEs:	Solve $f(\dot{x}(t), x(t), t) = 0$ $t \in [0, T], x(0) = x_0, \dot{x}(0) = x'_0$ for $x(t) \in \mathbb{R}^n, t \in [0, T]$		Rythmos
Nonlinear Problems	Given nonlinear operator $F(x) \in \mathbb{R}^m \rightarrow \mathbb{R}^m$ Solve $F(x) = 0 \quad x \in \mathbb{R}^n$		NOX
Linear Problems Linear Equations: Eigen Problems:	Given Linear Ops (Matrices) $A, B \in \mathbb{R}^{m \times n}$ Solve $Ax = b$ for $x \in \mathbb{R}^n$ Solve $A\nu = \lambda B\nu$ for (all) $\nu \in \mathbb{R}^n, \lambda \in \mathbb{C}$		Anasazi Ifpack, ML, etc... AztecOO
Distributed Linear Algebra Matrix/Graph Equations: Vector Problems:	Compute $y = Ax; A = A(G); A \in \mathbb{R}^{m \times n}, G \in \mathbb{S}^{m \times n}$ Compute $y = \alpha x + \beta w; \alpha = \langle x, y \rangle; x, y \in \mathbb{R}^n$		Epetra Teuchos



Solver Software Stack

Trilinos

Phase I packages

Phase II packages

Phase III packages: Manycore*, templated

Optimization Unconstrained: Constrained:	Find $u \in \mathbb{R}^n$ that minimizes $g(u)$ Find $x \in \mathbb{R}^m$ and $u \in \mathbb{R}^n$ that minimizes $g(x, u)$ s.t. $f(x, u) = 0$	Sensitivities (Automatic Differentiation: Sacado)	MOOCHO	
Bifurcation Analysis	Given nonlinear operator $F(x, u) \in \mathbb{R}^{n+m}$ For $F(x, u) = 0$ find space $u \in U \ni \frac{\partial F}{\partial x}$		LOCA	T-LOCA
Transient Problems DAEs/ODEs:	Solve $f(\dot{x}(t), x(t), t) = 0$ $t \in [0, T], x(0) = x_0, \dot{x}(0) = x'_0$ for $x(t) \in \mathbb{R}^n, t \in [0, T]$		Rythmos	
Nonlinear Problems	Given nonlinear operator $F(x) \in \mathbb{R}^m \rightarrow \mathbb{R}^m$ Solve $F(x) = 0 \quad x \in \mathbb{R}^n$		NOX	T-NOX
Linear Problems Linear Equations: Eigen Problems:	Given Linear Ops (Matrices) $A, B \in \mathbb{R}^{m \times n}$ Solve $Ax = b$ for $x \in \mathbb{R}^n$ Solve $A\nu = \lambda B\nu$ for (all) $\nu \in \mathbb{R}^n, \lambda \in \mathbb{C}$		Anasazi	
Distributed Linear Algebra Matrix/Graph Equations: Vector Problems:	Compute $y = Ax; A = A(G); A \in \mathbb{R}^{m \times n}, G \in \mathbb{S}^{m \times n}$ Compute $y = \alpha x + \beta w; \alpha = \langle x, y \rangle; x, y \in \mathbb{R}^n$		AztecOO Ifpack, ML, etc...	Belos* T-Ifpack*, T-ML*, etc.
			Epetra	Tpetra* Kokkos
			Teuchos	



Trilinos/Kokkos Node API



Generic Shared Memory Node

- Abstract inter-node comm provides DMP support.
- Need some way to **portably** handle SMP support.
- Goal: allow code, once written, to be run on **any parallel node**, regardless of architecture.
- **Difficulty #1**: Many different **memory architectures**
 - Node may have multiple, disjoint memory spaces.
 - Optimal performance may require special memory placement.
- **Difficulty #2**: **Kernels** must be tailored to architecture
 - Implementation of optimal kernel will vary between archs
 - No universal binary → need for separate compilation paths



Kokkos Node API

- Kokkos provides two main components:
 - Kokkos memory model addresses Difficulty #1
 - Allocation, deallocation and efficient access of memory
 - compute buffer: special memory used for parallel computation
 - New: Local Store Pointer and Buffer with size.
 - Kokkos compute model addresses Difficulty #2
 - Description of kernels for parallel execution on a node
 - Provides stubs for common parallel work constructs
 - Currently, parallel for loop and parallel reduce
- Code is developed around a polymorphic Node object.
- Supporting a new platform requires only the implementation of a new node type.



Kokkos Memory Model

- A generic node model must at least:
 - support the scenario involving **distinct device memory**
 - allow **efficient** memory access under traditional scenarios
- Nodes provide the following memory routines:

```
ArrayRCP<T> Node::allocBuffer<T>(size_t sz);  
void        Node::copyToBuffer<T>( T * src,  
                                   ArrayRCP<T> dest);  
void        Node::copyFromBuffer<T>(ArrayRCP<T> src,  
                                   T * dest);  
ArrayRCP<T> Node::viewBuffer<T> (ArrayRCP<T> buff);  
void        Node::readyBuffer<T>(ArrayRCP<T> buff);
```

Kokkos Compute Model

- How to make shared-memory programming generic:
 - **Parallel reduction** is the intersection of `dot()` and `norm1()`
 - **Parallel for loop** is the intersection of `axpy()` and mat-vec
 - We need a way of **fusing** kernels with these basic **constructs**.
- Template meta-programming is **the answer**.
 - This is the same approach that Intel TBB and Thrust take.
 - Has the effect of requiring that Tpetra objects be templated on Node type.
- Node provides generic parallel constructs, user fills in the rest:

```
template <class WDP>
void Node::parallel_for(
    int beg, int end, WDP workdata);
```

Work-data pair (WDP) struct provides:

- loop body via `WDP::execute(i)`

```
template <class WDP>
WDP::ReductionType Node::parallel_reduce(
    int beg, int end, WDP workdata);
```

Work-data pair (WDP) struct provides:

- reduction type `WDP::ReductionType`
- element generation via `WDP::generate(i)`
- reduction via `WDP::reduce(x, y)`

Example Kernels: axpy() and dot()

```
template <class WDP>
void
Node::parallel_for(int beg, int end,
                   WDP workdata    );
```

```
template <class T>
struct AxyOp {
    const T * x;
    T * y;
    T alpha, beta;
    void execute(int i)
    { y[i] = alpha*x[i] + beta*y[i]; }
};
```

```
AxyOp<double> op;
op.x = ...; op.alpha = ...;
op.y = ...; op.beta  = ...;
node.parallel_for< AxyOp<double> >
    (0, length, op);
```

```
template <class WDP>
WDP::ReductionType
Node::parallel_reduce(int beg, int end,
                     WDP workdata    );
```

```
template <class T>
struct DotOp {
    typedef T ReductionType;
    const T * x, * y;
    T identity()      { return (T)0;      }
    T generate(int i) { return x[i]*y[i]; }
    T reduce(T x, T y) { return x + y;    }
};
```

```
DotOp<float> op;
op.x = ...; op.y = ...;
float dot;
dot = node.parallel_reduce< DotOp<float> >
    (0, length, op);
```



Hybrid CPU/GPU Computing

Hybrid Timings (Tpetra)

- Tests of a simple iterations:
 - **power method**: one sparse mat-vec, two vector operations
 - **conjugate gradient**: one sparse mat-vec, five vector operations
- DNVs/x104 from UF Sparse Matrix Collection (100K rows, 9M entries)
- NCCS/ORNL **Lens** node includes:
 - one NVIDIA Tesla C1060
 - one NVIDIA 8800 GTX
 - Four AMD quad-core CPUs
- Results are **very tentative!**
 - suboptimal GPU traffic
 - bad format/kernel for GPU
 - bad data placement for threads

Node	PM (mflop/s)	CG (mflop/s)
Single thread	140	614
8800 GPU	1,172	1,222
Tesla GPU	1,475	1,531
Tesla + 8800	981	1,025
16 threads	816	1,376
1 node		
15 threads + Tesla	867	1,731
2 nodes		
15 threads + Tesla	1,677	2,102

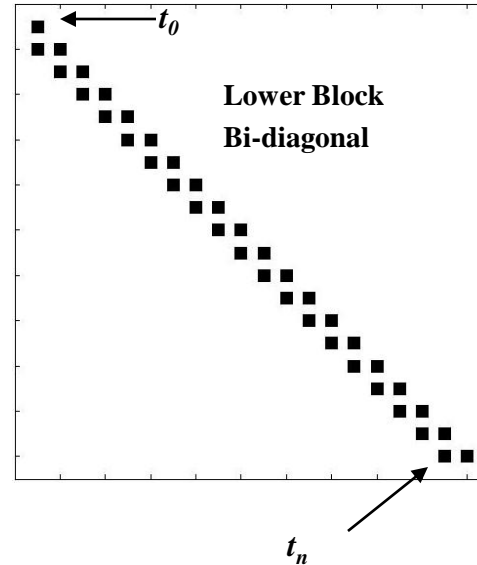
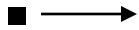
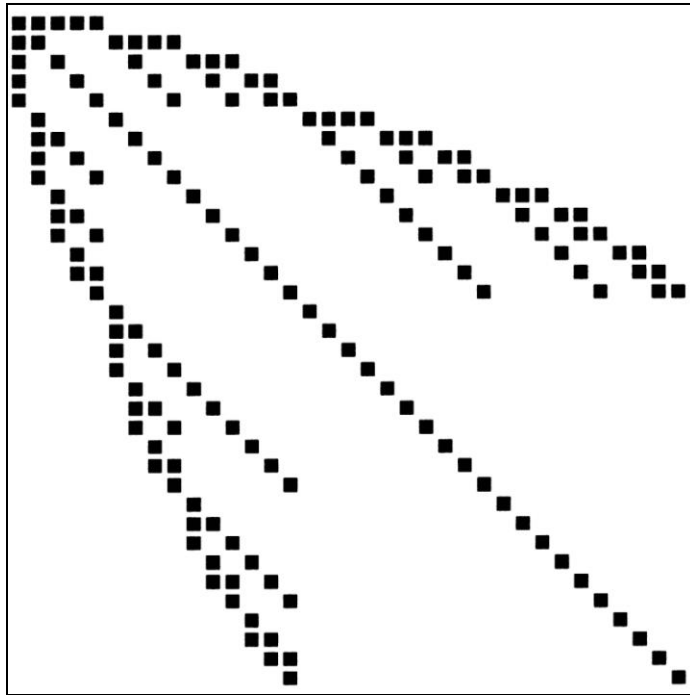


New Core Linear Algebra Needs

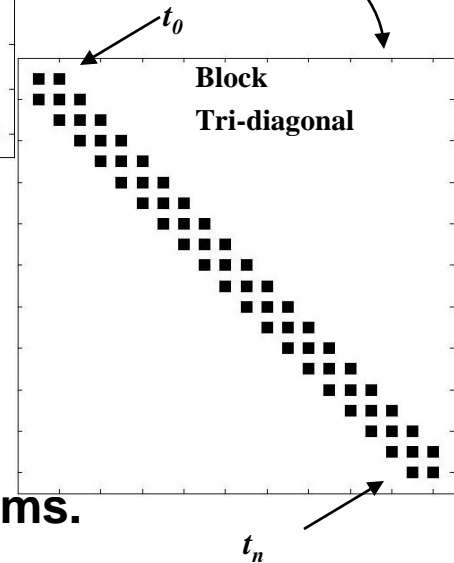
Advanced Modeling and Simulation Capabilities: Stability, Uncertainty and Optimization

- Promise: 10-1000 times increase in parallelism (or more).

SPDEs:



Transient
Optimization:



- Pre-requisite: High-fidelity “forward” solve:
 - Computing families of solutions to similar problems.
 - Differences in results must be meaningful.

■ - Size of a single forward problem



Advanced Capabilities: Readiness and Importance

Modeling Area	Sufficient Fidelity?	Other concerns	Advanced capabilities priority
Seismic <i>S. Collis, C. Ober</i>	Yes.	None as big.	Top.
Shock & Multiphysics (Alegra) <i>A. Robinson, C. Ober</i>	Yes, but some concerns.	Constitutive models, material responses maturity.	Secondary now. Non-intrusive most attractive.
Multiphysics (Charon) <i>J. Shadid</i>	Reacting flow w/ simple transport, device w/ drift diffusion, ...	Higher fidelity, more accurate multiphysics.	Emerging, not top.
Solid mechanics <i>K. Pierson</i>	Yes, but...	Better contact. Better timestepping. Failure modeling.	Not high for now.



Advanced Capabilities: Other issues

- Non-intrusive algorithms (e.g., Dakota):
 - Task level parallel:
 - A true peta/exa scale problem?
 - Needs a cluster of 1000 tera/peta scale nodes.
- Embedded/intrusive algorithms (e.g., Trilinos):
 - Cost of code refactoring:
 - Non-linear application becomes “subroutine”.
 - Disruptive, pervasive design changes.
- Forward problem fidelity:
 - Not uniformly available.
 - Smoothness issues.
 - Material responses.



Advanced Capabilities: Derived Requirements

- Large-scale problem presents collections of related subproblems with forward problem sizes.

- Linear Solvers: $Ax = b \rightarrow AX = B, Ax^i = b^i, A^i x^i = b^i$
 - Krylov methods for multiple RHS, related systems.

- Preconditioners:
$$A^i = A_0 + \Delta A^i$$
 - Preconditioners for related systems.

- Data structures/communication:
$$pattern(A^i) = pattern(A^j)$$
 - Substantial graph data reuse.



Summary

- App targets will change:
 - Advanced modeling and simulation: Gives a better answer.
 - Kernel set changes.
- Resilience requires an integrated strategy:
 - Most effort at the system/runtime level.
 - C/R (with localization) will continue at the app level.
 - Resilient algorithms will mitigate soft error impact.
- Building the next generation of parallel applications requires enabling domain scientists:
 - Write sophisticated methods.
 - Do so with serial fragments.
 - Fragments hoisted into scalable, resilient fragment.



Quiz (True or False)

1. MPI-only has the best parallel performance.
2. Future parallel applications will not have MPI_Init().
3. All future programmers will need to write parallel code.
4. Use of “markup”, e.g., OpenMP pragmas, is the least intrusive approach to parallelizing a code.
5. DRY is not possible across CPUs and GPUs
6. GPUs are a harbinger of CPU things to come.
7. Checkpoint/Restart will be sufficient for scalable resilience.
8. Resilience will be built into algorithms.
9. MPI-only and MPI+X can coexist in the same application.
10. Kernels will be different in the future.